Claims:

1) A switch fabric implemented on a chip, comprising:
   a) an array of cells;
   b) an I/O interface in communication with said array of cells for permitting exchange of data packets between said array of cells and components external to said array of cells;
   c) each cell communicating with at least one other cell of said array permitting exchange of data packets between the cells of said array;
   d) each cell including:
      I) a memory for receiving a data packet from another cell of said array;
      II) a control entity to control release of a data packet toward a selected destination cell of said array at least in part on a basis of a degree of occupancy of the memory in said destination cell.

2) A switch fabric as defined in claim 1, wherein each cell of said array includes:
   a) a transmitter in communication with said I/O interface and in communication with every other cell of said array, said transmitter operative to process a data packet received from said I/O interface to determine a destination of the data packet and forward the data packet to at least one cell of said array selected on a basis of the determined destination;
   b) a plurality of receivers associated with respective cells from said array, each receiver being in communication with a respective cell allowing the

respective cell to forward data packets to the receiver;

c) said receivers in communication with said I/O interface for releasing data packets to said I/O interface.

3) A switch fabric as defined in claim 2, wherein said array of cells includes a plurality of data channels, each data channel being associated with a given cell, the data channel associated with said given cell connecting the transmitter of said given cell to receivers in cells other than said given cell and associated with said given cell.

4) A switch fabric as defined in claim 2, wherein said array of cells includes a plurality of data channels, each data channel being associated with a given cell, the data channel associated with said given cell connecting the transmitter of said given cell to a receiver in every cell of said array of cells and associated with said given cell.

5) A switch fabric as defined in claim 4, wherein the plurality of data channels are independent from one another, wherein transmission of a data packet over one data channel is made independently of a transmission of a data packet over another data channel.

6) A switch fabric as defined in claim 5, wherein each data channel performs a parallel data transfer.

7) A switch fabric as defined in claim 1, wherein said array of cells forms a matrix.

8) A switch fabric as defined in claim 7, wherein said matrix is bi-dimensional.

9) A switch fabric as defined in claim 7, wherein said matrix is three-dimensional.

10) A switch fabric as defined in claim 1, wherein said array of cells forms a toroidal mesh arrangement.

11) A switch fabric as defined in claim 2, wherein said memory is a first memory and wherein the transmitter of said given cell includes a second memory for storing data packets received from said I/O interface.

12) A switch fabric as defined in claim 11, wherein said second memory includes a plurality of segments, each segment being associated with a receiver in a cell of said array to which the transmitter of said given cell is capable of forwarding a data packet via the data channel.

13) A switch fabric as defined in claim 12, wherein the transmitter of said given cell includes said control entity, said control entity being operative to process a data packet forwarded from said I/O interface to determine a cell of said array to which the data packet is destined and identify on a basis of the determined cell a segment of said second memory into which the packet is to be loaded.

14) A switch fabric as defined in claim 13, wherein said control entity includes a plurality of queue controllers associated with respective segments of said second memory.

15) A switch fabric as defined in claim 14, wherein said second memory implements a plurality of registers, each register being associated with a queue controller and being suitable for holding data representative of a degree of occupancy of a segment of said second memory associated with the queue controller.

16) A switch fabric as defined in claim 15, wherein a data packet received by said transmitter from said I/O interface is characterized by a priority level selected in a group of priority levels, each segment of said second memory being partitioned into slots, each slot capable of storing at least one data packet, each slot being associated with a given priority level of said group of priority levels.

17) A switch fabric as defined in claim 16, wherein the registers of said second memory associated with each queue controller store data indicative of a degree of occupancy of the slots of said segment associated with the queue controller, for each priority level of the group of priority levels.

18) A switch fabric as defined in claim 13, wherein said first memory is divided into a plurality of sectors associated with respective ones of said receivers, said

sectors capable of storing data packets forwarded to said receivers by cells of said array, said control entity being operative to communicate with each receiver associated with said given cell to assess a degree of occupancy of the sector of each receiver associated with said given cell.

19) A switch fabric as defined in claim 18, wherein said control entity communicates with each receiver associated with said given cell to assess the degree of occupancy of the sector of each receiver associated with said given cell, over a back channel.

20) A switch fabric as defined in claim 19, including a plurality of back channels, there being a dedicated back channel between said control entity and each receiver associated with said given cell.

21) A switch fabric as defined in claim 20, wherein each back channel transfers data serially.

22) A switch fabric as defined in claim 19, wherein said second memory includes an area for storing data indicative of the degree of occupancy of the sector of each receiver associated with said given cell.

23) A switch fabric as defined in claim 22, wherein said control entity is operative to process the data indicative of the degree of occupancy of the sector of each receiver associated with said given cell to determine which data packet stored in said second memory is suitable for transmission to a receiver.

24) A switch fabric as defined in claim 23, wherein when said control entity determines that a data packet is suitable for transmission, said control entity generates a control signal to request transmission of the data packet.

25) A switch fabric as defined in claim 24, wherein when said control entity determines that a plurality of data packets are suitable for transmission, said control entity generates a plurality of control signals to request transmission of the data packets, each control signal being associated with a data packet.

26) A switch fabric as defined in claim 25, wherein said control entity includes an arbiter for processing said control signals to select a data packet to transmit among the plurality of data packets suitable for transmission.

27) A switch fabric as defined in claim 26, wherein a data packet is characterized by a priority level, wherein each control signal conveys the priority level of the data packet associated with the control signal.

28) A switch fabric as defined in claim 27, wherein said arbiter selects a data packet to transmit among the plurality of data packets suitable for transmission on a basis of the priority levels of the plurality of data packets suitable for transmission.

29) A switch fabric as defined in claim 28 wherein said arbiter processes control signals to request transmission of data packets in a round robin manner.

30) A switch fabric as defined in claim 29, wherein said arbiter selects a data packet to transmit among the plurality of data packets suitable for transmission on a basis of the priority levels of the plurality of data packets suitable for transmission and on the basis of whether or not a data packet was previously submitted for transmission.

31) A switch fabric as defined in claim 18, wherein each receiver of said plurality of receivers communicates with said I/O interface.

32) A switch fabric as defined in claim 31, wherein said control entity is a first control entity and wherein said plurality of receivers include a second control entity to regulate a release of data packets from said sectors to said I/O interface.

33) A switch fabric as defined in claim 32, wherein said second control entity includes a plurality of queue controllers associated with respective sectors of said first memory.

34) A switch fabric as defined in claim 33, wherein a data packet received by a receiver of said plurality of receivers is characterized by a priority level selected in a group of priority levels, each sector of said second memory being divided into subdivisions, each

subdivision capable of storing at least one data packet, each subdivision being associated with a given priority level of said group of priority levels.

35) A switch fabric as defined in claim 34, wherein said second control entity includes an arbiter in communication with said queue controllers, each queue controller operative to transmit a control signal to said arbiter for each data packet held in the sector associated with the queue control to request release of the data packet to said I/O interface.

36) A switch fabric as defined in claim 35, wherein each control signal conveys the priority level of the data packet associated with the control signal.

37) A switch fabric as defined in claim 36, wherein said arbiter selects a data packet for release to said I/O interface among the data packets corresponding to the control signals transmitted to said arbiter on the basis of the levels of priority of the data packets corresponding to the control signals transmitted to said arbiter.

38) A switch fabric as defined in claim 2, wherein each data packet comprises a plurality of words including a first word of said data packet and a last word of said data packet, wherein each word comprises a field indicative of whether said word is a pre-determined number of words away from said last word of said data packet.

39) A switch fabric as defined in claim 38, wherein the transmitter is operative to monitor said field in each word of each data packet forwarded to at least one cell of said array, the transmitter further being operative to begin forwarding a next data packet upon detecting that said field of a word in a packet currently being forwarded is indicative of said word being a pre-determined number of words away from the last word of said data packet currently being forwarded.

40) A switch fabric as defined in claim 2, each cell further including a central processing unit (CPU) connected to the transmitter, said transmitter being further operative to process a data packet received from said CPU to determine a destination of the data packet and forward the data packet to at least one cell of said array selected on the basis of the determined destination.

41) A switch fabric as defined in claim 3, each cell further including a central processing unit (CPU) connected to the transmitter, said transmitter being further operative to process a data packet received from said CPU to determine a destination of the data packet and forward the data packet to at least one cell of said array selected on the basis of the determined destination, wherein data packets received by the transmitter in a given cell from the I/O interface and from the CPU in said given cell share the data channel associated with said given cell.

42) A switch fabric as defined in claim 2, each cell further including a central processing unit (CPU) connected to the plurality of receivers, said receivers being further operative to determine whether data packets are to be released to the I/O interface or to the CPU and release said data packets accordingly.

43) A switch fabric as claimed in claim 42, wherein each data packet comprises a field indicative of whether the data packet is destined for a CPU and wherein said receivers are operative to determine whether data packets are to be released to the I/O interface or to the CPU on the basis of said field.

44) A switch fabric as defined in claim 25, each cell further including a central processing unit (CPU) connected to the plurality of receivers, wherein said control entity includes a first arbiter for processing said control signals to select a data packet to transmit to the I/O interface among the plurality of data packets suitable for transmission to the I/O interface, wherein said control entity includes a second arbiter for processing said control signals to select a data packet to transmit to the CPU among the plurality of data packets suitable for transmission to the CPU.